

The Geographic Flow of Music

Conrad Lee and Pádraig Cunningham

Clique Research Cluster

University College Dublin

8 Belfield Office Park, Clonskeagh

Dublin 4, Ireland

Tel: +353 1 716 5346

Email: conradlee@gmail.com, padraig.cunningham@ucd.ie

Abstract—The social media website last.fm provides a detailed snapshot of what its users in hundreds of cities listen to each week. After suitably normalizing this data, we use it to test three hypotheses related to the geographic flow of music. The first is that although many of the most popular artists are listened to around the world, music preferences are closely related to nationality, language, and geographic location. We find support for this hypothesis, with a couple of minor, yet interesting, exceptions. Our second hypothesis is that some cities are consistently early adopters of new music (and early to snub stale music). To test this hypothesis, we adapt a method previously used to detect the leadership networks present in flocks of birds. We find empirical support for the claim that a similar leadership network exists among cities, and this finding is the main contribution of the paper. Finally, we test the hypothesis that large cities tend to be ahead of smaller cities—we find only weak support for this hypothesis.

I. INTRODUCTION

The question of how information and preferences spread through social networks has a long and rich history. The topic became an active field of study just after World War Two [1]–[4]; this early work produced, for example, the seminal *two-step flow of communication* hypothesis, which states that “ideas often flow from radio and print to opinion leaders, and from these to the less active sections of the population” [5]. In the 1970s, Mark Granovetter contributed prominent ideas to the field, including the hypothesis that members of tightly-knit social groups have largely duplicate information, and rely on acquaintanceships with members of other groups to gain access to novel information [6].

More recently, detailed logs of digital communication have enabled these hypotheses to be tested on datasets that are much larger than was feasible only a decade earlier. For example, in [7], Bakshy et al. subject 250 million facebook users to a controlled experiment in order to measure the role that facebook friends play in influencing the diffusion of information, finding that while a user’s most active relationships are individually the most influential, the overall effect of less active relationships in spreading novel information is stronger. Additional examples of recent significant work includes the worldwide spread of e-mail chain letters [8], the analysis of a massive worldwide instant messaging dataset [9], and the spread of information through the blogosphere [10], [11].

Here, we investigate hypotheses related to the geographic flow of preferences in music. Our main contribution is to for-

malize and answer the following question: if one considers the month-by-month change in the aggregate musical preferences of cities, are some cities consistently ahead of others? In other words, can we find that some cities are leaders and others are followers?

Our enquiry into the geographic distribution of musical preferences is structured as follows. We begin by describing the data, a world-wide log of listening habits recorded by last.fm, as well as various pre-processing and normalization steps in section II. Next, in section III, we measure how regional musical preferences are, finding that although many of the most popular artists are popular all around the world, there are nonetheless well-defined clusters of cities that are closely related to nationality, language, and geographic distance.

In section IV, we move on to our main contribution: an analysis of the *dynamics* of music preferences. We adapt a methodology previously used to find leadership in pigeon flocks [12] to detect whether some cities consistently follow others. At a high level, this methodology involves looking at every dyad (pair of nodes) and running a test to see whether the time-lagged correlation is larger in one direction than another. We observe that when we put all of these directed pairs together, the resulting networks are nearly acyclic, a strong indicator that the geographic flow of music has a clear direction, i.e., hierarchical structure [13].

Recently there has been much excitement surrounding the observation that productivity, efficiency, and innovation all scale super-linearly with the size of a city [14]–[16]. This line of reasoning suggests the hypothesis that larger cities should also be more up to date on the latest and greatest music. We wrap up our inquiry into the spread of music in section VI by testing the hypothesis that leadership is also correlated with the size of a city.

II. DATA: PREPROCESSING & NORMALIZATION

Last.fm is a service based around collecting data on the listening habits of its users. Users install a plug-in on their audio players such as iTunes or Winamp which keeps track of the songs that the user listens to, either on his computer or external device (e.g., an iPod). The plug-in uploads this information to the last.fm database, giving the service a log of what its users listen to. In 2011 alone, last.fm received 11 billion such notifications (called “scrobbles” by last.fm), and

since the service began in 2003 it has received 61 billion.¹ Last.fm uses this information in various ways, for example, to compare the similarity of two users' musical taste, to recommend music, and to create a profile page.

Creating listen matrices. Last.fm aggregates this data into weekly charts for over 200 metropolitan areas around the world, and makes the data behind these charts accessible through a public API. For every week and each city, the last.fm API indicates the number of unique listeners that each of that city's top 500 artists had. Thus, for each week we have a matrix; in this matrix every city is a row vector with 500 non-zero elements, and each column represents an artist. Because not all cities have the same top 500 artists, the matrix has more than 500 columns and a large number of zero-valued elements. Thus, a non-zero entry in this matrix at position i, j is a positive integer indicating the number of unique users from city i who listened to artist j that week. Zero-valued entries indicate that the artist had either no listeners, or that it was not among the 500 most popular artists in the city that week. At the time of data collection in late 2011, these charts were available for 153 weeks.

Because not all last.fm users are active every week, a single week's chart can be thought of as a sample of listening preferences among last.fm users. In cities that have relatively few users, the variance associated with this sample becomes large, indicating noise. We find that we can reduce this noise by summing up the matrices associated with four consecutive weeks together. This effectively increases the sample size for each entry in the city-artist matrix described above. For this reason, in all of the analysis below, we aggregate our data using a "sliding window" where the width of the window is four weeks, and the window slides in one-week steps. We call the matrices associated with these four-week periods *listen matrices*.

Normalizing listen matrices. Consider the toy example presented in fig. 2(A). In this scenario, we imagine there are only two artists, Radiohead and Coldplay, and two cities, Los Angeles and Seattle. We want to compare how similar Los Angeles' preferences are to Seattle's. In one sense, they are similar: each city listens to roughly 50% more Radiohead than Coldplay. However, if we look at the absolute number of listens in each city, the cities are far apart simply because Los Angeles is much larger than Seattle.

In order to compare the similarities of cities regardless of their size (i.e., last.fm activity level), we always perform Euclidean normalization on the rows of each listen matrix, which ensures that each row vector (i.e., each city's listening preference) has the same length. In other words, the Euclidean normalization puts the row vectors of each listen matrix on the unit circle, as in fig. 2(B). This type of normalization is standard in the field of Information Retrieval [17].

Genres. In the analysis below, it will be important to

distinguish between various genres of music. In order to determine which genres exist, and which artists belong to each genre, we use last.fm's *tag* API. Examples of tags include rock, seen live, alternative, indie, electronic, and pop (these are the 6 most popular tags). For each tag, the last.fm API also indicates the one thousand most popular artists that belong to that tag. We construct the listen matrix associated with a given tag by including only those columns which represent artists included in the list of top thousand artists for that tag. We will subsequently refer to the term "tag" by the more conventional term "genre". Some tags, e.g. "seen live" are clearly not genres - these are not considered in the analysis presented here.

Missing data. Inspection of the data indicates that fourteen of the weeks are outliers in the sense that around the world, little if any music was listened to. We believe that during these weeks the last.fm scrobbling service was not operating as usual. In the analysis below, we omitted from all measurements the contributions that involved one or more of the missing weeks.

III. MUSIC KNOWS NO BORDERS, YET GEOGRAPHIC CLUSTERS ARE STRONG

Are the listening preferences of last.fm users across the world similar, or do they form coherent clusters? The existence of global superstars might lead one to believe that largely similar music is listened to across the world; indeed, in a comprehensive study of the top-40 music charts of 22 countries, Ferreira and Waldfogel found that 31 artists appeared simultaneously on at least 18 countries in one year [18]. Of these 31 artists, 23 were US American. That such a small set of artists appeared on charts all around the world suggests a high degree of homogeneity around the world.

Despite this appearance of global homogeneity, in this section, we present results which indicate that there are clusters of cities that have their own idiosyncratic preferences, and that these clusters are closely related to geographic distance, nationality, and language.

Producing a hierarchical clustering. To construct the dendrogram shown in fig. 1, we performed average linkage clustering (an agglomerative clustering algorithm) on a distance matrix D of the cities, a square matrix where each entry $D_{i,j}$ is the Euclidean distance between city i and j . Instead of constructing the dendrogram based on just a single listen matrix, we summed together the distance matrices associated with all of the listen matrices in our dataset. The colored clusters are the result of taking a flat cut to the dendrogram at a height which we chose manually. For an overview of this type of hierarchical clustering, as well as a description of the software package we used, see [19].

Discussion. If we look at the lowest level structure of the dendrogram—i.e., the pairs of cities that are most similar to each other—we observe that every pair involves two cities of the same nationality. Many of these pairs are composed of cities that are, in the context of their countries, geographically close

¹According to last.fm's blog post at <http://blog.last.fm/2012/01/16/building-best-of-2011>




Fig. 1. Hierarchical clustering based on average linkage clustering of Euclidean distances of cities in the normalized listen matrices. (Zoom in electronic version to view city names.)

to each other: Cincinnati and Columbus, Portland and Seattle, Berlin and Dresden, Edmonton and Calgary, Lausanne and Geneva, Gijon and Oviedo, Birmingham and Manchester, and Edinburgh and Glasgow. However, there are a few noteworthy exceptions to this trend: New York City and San Francisco, Milan and Rome, Munich and Hamburg. These pairs of cities are close in the dendrogram indicating high relative similarity. This is surprising because in each case both cities in the pair have many geographically nearer cities which would seem to be more likely candidates for most similar counterpart. For example, San Jose is geographically adjacent to San Francisco, but San Francisco's users are more similar to NYC's.

At an intermediate level of structure in the dendrogram—the colored clusters—we see that again, nationality dominates. In the cases where two countries have been put into the same cluster, such as New Zealand and Australia or Ireland and the United Kingdom, those two countries still show separation within their cluster. Intriguingly, the United States shows two somewhat distinct clusters which are not geographically coherent and are difficult to explain. We note only that the USA 2 cluster appears to contain both the largest metropolitan areas (NYC, Los Angeles, and Chicago), as well as several cities known for having a large “hipster” population, which is passionate about appearing to know much about music, and likely to use last.fm (such as San Francisco, Austin, Portland, and Seattle).

At the highest level of the hierarchical structure, we observe how the colored clusters are related to each other. Language seems to be the key here; Anglophone clusters are closely related, as are the Spanish-speaking and German-speaking clusters. It is interesting to see that Switzerland, including its French-speaking cities, are most closely related to German speaking countries. Chilean cities' musical preferences are more similar to Mexican cities than to Brazilian cities. We note that although in this dendrogram Canada appears to be more similar to N. Zealand and Australia than to the USA, this tendency was not very robust: if we ran the clustering on a subset of only more active cities, then it was usually more similar to the US cities. On the other hand, the other features we have described here were robust in this sense.

IV. METHODOLOGY: DETECTING LEADERS AND FOLLOWERS

To detect leader-follower pairs, we adapt the methodology of Nagy et al. [12], which is based on finding lagged correlations, and was previously applied to finding leadership in pigeon flocks. In fig. 2, we display some of the key steps of the method we employ to find leaders and followers. Here we show made-up data for explanation purposes: we depict a scenario with just two cities, Los Angeles (LA) and Seattle, and two artists. We are interested in determining whether

- LA follows Seattle (in this case we draw the directed edge $LA \rightarrow Seattle$)
- Seattle follows LA (Seattle $\rightarrow LA$), or
- neither leads the other (no edge)

If an edge exists, we would also like to assign a weight to that edge which determines the strength of the leader-follower relationship. We now explain how we decide on the relationship type and weight.

Calculating lagged correlations. We begin by performing Euclidean normalization on each city’s listening frequency vector in every listen matrix, as previously described in section II and visualized in the change from fig. 2(A) to fig. 2(B). Each of the blue arrows in fig. 2(B) is a velocity $v_{\text{seattle}}(t, t + 1)$ that represents the change that takes place in the listening habits of Seattle from one month t to the next month $t + 1$. For example, to find Seattle’s velocity from June to July $v_{\text{seattle}}(\text{June}, \text{July})$, we subtract Seattle’s row in the normalized listen matrix for time-step June from the corresponding row from the matrix for July.

As mentioned in section II, each listen matrix is based on a four-week window of last.fm data, which means that to calculate one of these velocities, we use eight consecutive weeks (two four-week windows). We successively slide this eight week period one week forward in time, giving us one velocity associated with each slide. We are left with a sequence of velocities for each city.

To measure whether Seattle follows LA, we measure the similarity of each of Seattle’s velocities with LA’s velocities from one month earlier, as in the top half of fig. 2(C). We measure the similarity between two velocities using the dot product (as in [12]). We call the average of these lagged similarities the *correlation of LA’s velocities with Seattle’s lagged velocities*, where the lag size is one month, and we refer to this measure as C .

In the example displayed in fig. 2, the lag size is fixed at one month. However, there is no reason to believe that this lag size should be the same for all dyads and it would be arbitrary to settle on one month. Along the lines of Nagy et al., for each dyad, we consider lag sizes of 1-5 weeks, and we choose the one which yields the largest correlation. We therefore let the data decide how this parameter ought to be set. In practice the lag size which maximizes the correlation tends to be one week, however there are also cases where the strongest correlation is at four or five weeks (see blue edges in fig. 3) – in these cases the correlation tends to be weak.

Deciding which edges to accept. Up to now, the methodology described in this section closely resembles the one used by Nagy et al. However, we find it necessary to modify their final two steps, which determine

- (1) whether a correlation is strong enough to be accepted
- (2) the direction the relationship if one exists.

For step (1), Nagy et al. accept only those leader-follower relationships which have a correlation above some threshold, either 0.5 or 0.9. This criterion is inappropriate in our case because the magnitude of the dot products are very small, on the order of 0.01 to 0.001. They are much smaller because, due to the way we normalize data, cities mostly stand still and move only slightly from week to week; furthermore we are in a much higher dimensional space. For these reasons it is hard

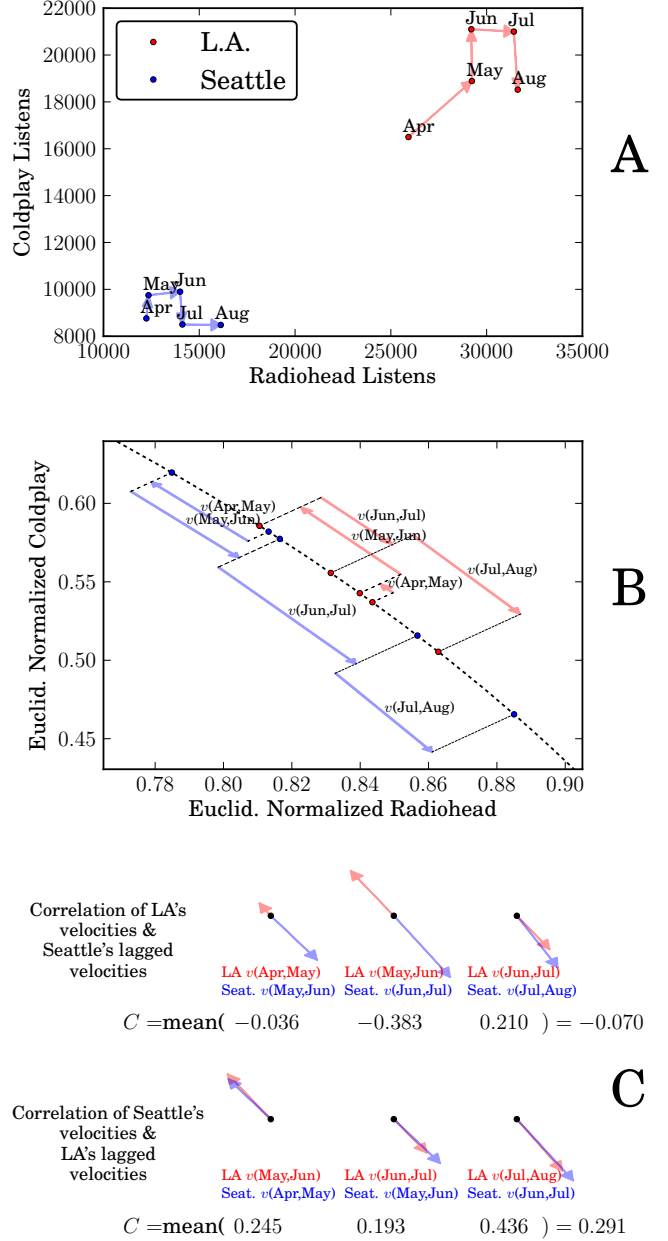


Fig. 2. **Calculating lagged correlations** (here shown with imaginary data for ease of explanation). (A) First, for each city, we collect from last.fm the number of times that each artist was listened to in a given month. (B) To be able to compare cities with different levels of last.fm activity, we next normalize the number of listens in each city by that city’s Euclidean norm. We focus on the velocity (change in the normalized artist popularity) from the previous month $i - 1$ to the current month i , denoted as $v_j(i - 1, i)$ for city j , and depicted by the arrows in (B). (C) For each pair of cities (j, k) , we measure the similarity of $v_j(i - 1, i)$ and $v_k(i - 2, i - 1)$ by taking the dot product of these velocities. This yields a list of similarities over time; we define the lagged correlation to be the mean of these dot products. In this toy example, it should be clear from glancing at the trajectories of Seattle and LA that LA is following Seattle, and not the other way; the correlation measure presented here successfully indicates this tendency.

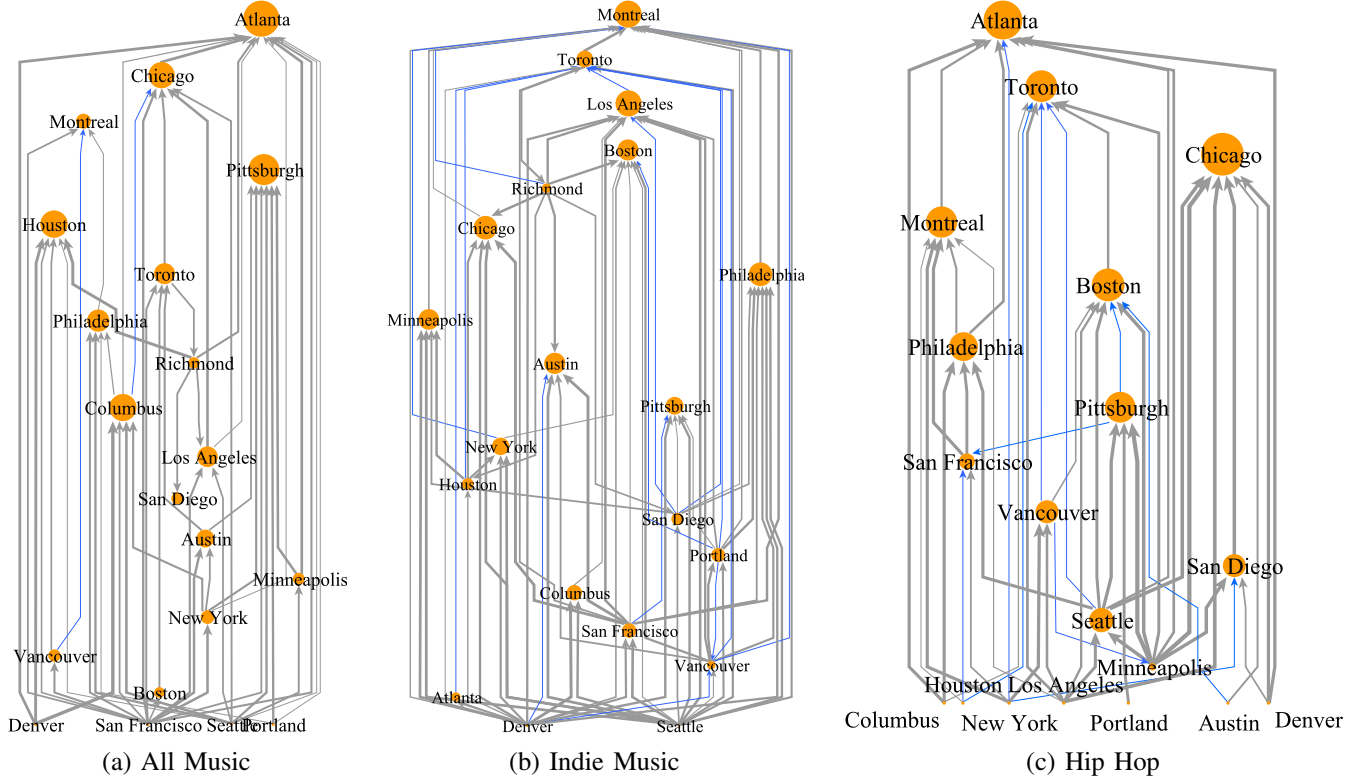


Fig. 3. Leader-follower network for the twenty most active cities in two Canada and the USA. Within each diagram, the height of the nodes corresponds to their PageRank, their size to their weighted in-degree. Edge width is determined by the correlation, as defined in section IV. Gray edges have a lag time of 1, 2, or 3 weeks, blue edges have a lag time of 4 or 5 weeks.

to pick a threshold for the lowest admissible correlation size. Instead, we perform one sample t -test on the distribution of dot products. If we cannot reject the hypothesis that the mean of the distribution equals zero, then we say no leader-follower relationship exists.

It could be the case for a dyad i, j that after performing step (1), i appears to follow j and j appears to follow i . While in this case Nagy et al. simply choose the direction that is larger (even if it is just marginally larger), we argue that in this situation perhaps neither city is really leading the other, and instead they are moving together. To make sure there is a clear direction to the leader-follower relationship, we perform a second t -test to make sure that the two correlations (which are means of dot products) are not equal; here we use a two-sided, paired t -test. If a one correlation is larger, then we accept the leader-follower pair associated with that correlation as a directed edge, otherwise we say no leader-follower relationship exists.

In the following results, we set $p = 0.01$ for all t -tests. We note that our use of t -tests here is heuristic; for example, we do not test to make sure that the distribution of dot products is Gaussian (although they do appear reasonably symmetric and we obtained qualitatively the same results when outliers were removed), and we do not correct for our testing of multiple hypotheses. We use the t -tests as a selection criterion to identify the more pronounced leader-follower relationships,

not because we rely on their validity in a statistical sense.

V. RESULTS: THE GEOGRAPHIC FLOW OF MUSIC

In the previous section, we described how we determine whether a leader-follower relationship exists between two nodes. In each study displayed in fig. 3, we take a subset of cities, find all follower-relationships among them, and plot the resulting network. The edges point from followers to leaders and are weighted by the lagged correlation, as defined above.

To create the networks in fig. 3, we first choose a genre of music. While it is possible to create a network showing the flow of all genres of music, as we have done in fig. 3a and fig. 4a, we find this has a disadvantage: depending upon the genre that one considers, contradictory relationships may exist. For example, if we consider hip hop music as in fig. 3c, then we see that Atlanta has the most prominent position, whereas if we consider indie music as in fig. 3b, Atlanta has one of the least prominent positions. By considering all genres at once as in fig. 3a, these trends get washed out by the multi-dimensional aspect that genre brings to the data.

In fig. 3, we show the leader-follower relationships between 20 cities in the USA and Canada with the largest number of active last.fm users. We choose this subset because of the noise associated with small cities that have insufficient data, and because due to space constraints it is hard to visualize large networks. The most significant property visible in these




Fig. 4. Leader-follower network for the most active cities in Western Europe. Sizes, positions, and colors as in fig. 3. Zoom-able online.

TABLE I
LEADER-FOLLOWER NETWORKS HAVE FEW CYCLES

Region	Genre	% Edge weight removed to make acyclic
N. America	All	0.0%
	Indie	1.8%
	Hip hop	0.0%
	Rock	0.0%
	Classic Rock	0.0%
Europe	All	0.0%
	Indie	0.0%
	Hip hop	2.2%
	Rock	0.0%
	Classic Rock	0.0%

graphs is that they are nearly acyclic; for example, fig. 3a has no edges, and by removing only three edges from fig. 3b, the graph becomes acyclic. In table I, we show that this property holds true for the leader-follower networks created from diverse geographic regions and genres. To calculate the measure displayed in that table, we first computed the feedback arc (edge) set, which is the smallest set of edges that, when removed from a graph, make the graph acyclic. We measured the percent of the graph’s total weight in the feedback arc set.

We believe that this lack of cycles is not an artifact of our methodology, which focuses only on dyads and does not consider the network as a whole. Rather, we believe that the

lack of cycles is inherent in the data itself, indicating a clear direction in the flow of music preferences. Others have argued that a system with a strong leadership hierarchy ought to be nearly acyclic [12], [13] so the lack of cycles in our networks is a clear validation of the methodology.

There are many centrality measures that could be used as criteria for deciding which cities are the most cutting edge and which are laggards. The networks in figs. 3 and 4 display two of these centrality measures: their height reflects their PageRank, which seems appropriate because PageRank is designed to rank importance of nodes on weighted, directed networks on which a dynamic process takes place [20]. The area taken up by each node reflects its weighted in-degree. While it is apparent that the PageRank and weighted in-degree are highly correlated, in some cases they order nodes differently—for example, in fig. 3c, Atlanta has the largest PageRank, but Chicago has the largest weighted in-degree. These visualizations were created using the “status” layout algorithm of the network visualization software Visone [21].

For us, the most surprising features of fig. 3 are (1) the middle ranking positions of some of the largest cities, such as NYC and LA in fig. 3a and NYC and Chicago in fig. 3b and (2) the prominent position of Canadian cities, especially in fig. 3b. While Montreal is known for having produced some popular indie bands (such as Arcade Fire and Wolf Parade), this does not necessarily mean that last.fm listeners from Montreal would be generally leaders in their taste in indie music; in any case, New York City is presumably home

to more prominent indie artists than Montreal.

While the diagrams in fig. 3 display the leader-follower relations for a relatively homogeneous cultural region, those in fig. 4 display these relations in Europe, a region more culturally and linguistically diverse. It is interesting to note that many of the most heavily weighted edges are between cities in different countries and which speak different languages. For example, London, Birmingham, Brighton, and Bristol, have a much stronger follower relationship with Oslo and Stockholm than with each other (London’s unremarkable position is also noteworthy). Similarly, Cracow and Warsaw do not follow each other, rather their strongest edges point to German and Scandinavian cities.

Along the lines of this last observation, it is noteworthy that in general many of the edges with the largest weights connect cities which were not similar to each other in the hierarchical clustering in fig. 1. For example, the Canadian cities are located far away from the US cities in that clustering, yet here there is a strong flow from the former to the latter. Although pairs of cities such as Portland and or NYC and San Francisco are very similar in the clustering, they are connected in fig. 3b by only weak edges. One speculative explanation is that cities which have very similar listening habits are largely synchronized with each other, and therefore there is little potential for novel information to flow between them. For example, the leading city in fig. 3b, Montreal, is unique in that the language spoken by the majority is not English but French, a difference which may provide it with novel information.

VI. HYPOTHESIS: LARGE CITIES ARE LEADERS

As noted in the introduction, there is currently much excitement surrounding the observation that productivity, efficiency, and innovation all scale super-linearly with the size of a city. For an accessible, high-level overview of this discussion, see [14]; for extensive empirical evidence for the universality of this relationship, see [15]; and for a proposed causal mechanism, see [16].

This work makes many fascinating empirical observations as well as an interesting comparison between organisms and cities; here we summarize only a few main points. The first is that the total productivity of a city P is super-linear. In data collected so far, a power-law relationship appears to provide a reasonable fit, so that total production in an N person city is well approximated by the relationship $P(N) = N^\beta$, where $\beta \approx 1.2$. [15] This means, for example, that a person living in cities with 10 million inhabitants is roughly 2.5 times as productive (in terms of wealth production, creativity, patents, and other measures) as an individual living in a city with only 100 thousand inhabitants. Consumption of water, gasoline, or electricity appear to have a linear relationship, so people in smaller and larger cities consume the same amounts. Certain types of infrastructural needs, such as the number of gasoline stations, the meters of electric cabling installed, and road surface area, increase sub-linearly, with the scaling exponent $\beta \approx 0.8$, indicating economies of scale.

TABLE II
RELATIONSHIP BETWEEN A CITY’S POPULATION SIZE AND STATUS IN NETWORK

Genre	Spearman rank correlation of centrality & population		% Edge weight where leader larger
	PageRank	In-degree	
All	0.34	0.18	55%
Indie	0.61	0.61	61%
Rock	0.21	0.26	63%
Hip Hop	0.38	0.28	59%

Bettencourt et al., the authors of [15], also suggest that the very pace of life in large cities is faster, and Arbesman et al. [16] propose that productivity gains may be attributed to the increased probability of ties between diverse groups, which helps information spread quickly. If the pace of life in larger cities were faster, and the spread of information more efficient, then it would be reasonable to expect that larger cities would lead smaller ones in adopting fresh music and abandoning stale music. Here we test this hypothesis by measuring whether city size is positively and strongly correlated with a position of leadership in the network flow diagrams presented in section V.

Bettencourt et al. are careful to treat each “national urban system” separately, because otherwise their measurements might be confounded by the fact that different countries have economies at different levels of development. Thus, they do not expect that all cities around the globe which are of the same size should have the same level of production; rather, they expect this only within a tightly integrated economic region. (They do however argue that the same scaling exponent exists in every nation.) The North American cities in fig. 3 belong to a tightly integrated economic area at a similar level of development, so we test this hypothesis on that set of cities. For US population sizes, we use the US Metropolitan Statistical Areas, as Bettencourt et al. (although we use the newer data from 2010), for Canadian population sizes we use the Census Metropolitan Areas from 2011.

In table II we display some measures of the relationship between the population size of a city and its leadership status in the diagrams depicted in fig. 3. The second and third columns display the Spearman’s rank correlation of population size and PageRank, and population size and the weighted in-degree, respectively. The final column shows, for all edges, the percentage of the total weight that comes from edges where the larger city is the leader.

While these correlations between city size and leadership position are positive, most of these relationships are quite weak when compared with those observed in the above-mentioned work on superlinearity of cities. We were surprised that they were not stronger. In most genres, the Spearman correlation coefficients are smaller than we expected, and the percentage of edge weight that comes from edges where small cities are led by larger cities is not very far from 50%. Additionally, although there are some cities in North America which dwarf

most of the other large cities (such as NYC, with 18.9 million residents, and LA with 12.8 million), in many cases these cities do not occupy prominent positions in fig. 3.

Indie music is an exception—here the correlation between leadership and city size is quite large. We are not sure why this is the case—perhaps this genre is quicker moving or more urban than the others (although presumably hip hop is also quite an urban genre).

The work on scaling laws in cities which we have summarized in this section is significant because it appears to have uncovered a universal law in the social sciences, one which can make quantitative predictions. Our point here is not to claim that our results contradict this law. The preliminary results presented in this paper suggest that, in the specific context of being cutting edge in music, cities are idiosyncratic. Larger cities are not predictably and generally ahead of smaller cities. In other words, a city is more than the number of its inhabitants, it might lead the trends in one genre while lagging them in another.

VII. DISCUSSION AND FUTURE WORK

One major question hangs over the results presented above: why should we believe that our models of flow, as pictured in the network diagrams displayed in this paper, are valid? On the one hand, two aspects of our methodology lend the results credibility: that each leader-follower relationship underwent a t -test, and that when all of the leader-follower relationships were put together into a graph, they formed directed acyclic graphs, which indicate a direction of flow in a strict sense. For these two reasons, our method distinguishes itself from other unsupervised methods—such as many clustering methods—which are problematic because they return results regardless of whether there is structure in the underlying data. In other words, if we shuffle our data around so that random noise dominates any signal of leader-follower relationships, our method no longer detects leader-follower relationships.

On the other hand, certain doubts remain, and we should stress that our results reflect a work in progress. For example, a relationship can be statistically significant but at the same time have a very small magnitude. We would be more confident of our results if we could demonstrate that the model that we create is meaningfully predictive. That is, given our model of leader-follower relationships among cities, and given a record of past listening behavior, we should be able to predict the changes in listening behavior that will occur in the near-term future better than a reasonable baseline predictor. We have not yet demonstrated that our models have this predictive power, although we plan to attempt this validation in future work.

ACKNOWLEDGMENT

This material is based upon works supported by the Science Foundation Ireland under Grant No. 08/SRC/I1407: Clique: Graph & Network Analysis Cluster.

REFERENCES

- [1] P. Lazarsfeld, B. Berelson, and H. Gaudet, *The People's Choice. How the Voter Makes up His Mind in a Presidential Campaign [1944]*. Columbia University Press, 1948.
- [2] J. Coleman, E. Katz, and H. Menzel, "The diffusion of an innovation among physicians," *Sociometry*, vol. 20, no. 4, pp. 253–270, 1957. [Online]. Available: <http://www.jstor.org/stable/2785979>
- [3] J. Moreno, "Who shall survive?, rev. ed.," *Beacon, NY: Beacon House*, 1953.
- [4] L. Festinger, S. Schachter, and K. Back, *Social pressures in informal groups: a study of human factors in housing*. Harper, 1950.
- [5] E. Katz, "The Two-Step Flow of Communication: An Up-To-Date Report on an Hypothesis," *The Public Opinion Quarterly*, vol. 21, no. 1, pp. 61 – 78, 1957. [Online]. Available: <http://www.jstor.org/stable/2746790>
- [6] M. Granovetter, "The strength of weak ties," *American journal of sociology*, pp. 1360–1380, 1973.
- [7] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic, "The role of social networks in information diffusion," *Arxiv preprint arXiv:1201.4145*, 2012.
- [8] D. Liben-Nowell and J. Kleinberg, "Tracing information flow on a global scale using Internet chain-letter data," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 105, no. 12, pp. 4633–8, Mar. 2008.
- [9] J. Leskovec and E. Horvitz, "Planetary-scale views on a large instant-messaging network," in *Proceeding of the 17th international conference on World Wide Web*. ACM, 2008, pp. 915–924. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1367497.1367620>
- [10] D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins, "Information diffusion through blogspace," in *Proceedings of the 13th international conference on World Wide Web*. ACM, 2004, pp. 491–501.
- [11] J. Leskovec, L. Backstrom, and J. Kleinberg, "Meme-tracking and the Dynamics of the News Cycle," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM New York, NY, USA, 2009, pp. 497–506.
- [12] M. Nagy, Z. Akos, D. Biro, and T. Vicsek, "Hierarchical group dynamics in pigeon flocks," *Nature*, vol. 464, no. 7290, pp. 890–3, Apr. 2010. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/20376149>
- [13] E. Mones, L. Vicsek, and T. Vicsek, "Hierarchy measure for complex networks," *Arxiv preprint arXiv:1202.0191*, 2012.
- [14] L. Bettencourt and G. West, "Bigger cities do more with less," *Scientific American*, vol. 305, no. 3, pp. 52–53, 2011.
- [15] L. Bettencourt, J. Lobo, D. Helbing, C. Kühnert, and G. West, "Growth, innovation, scaling, and the pace of life in cities," *Proceedings of the National Academy of Sciences*, vol. 104, no. 17, p. 7301, 2007.
- [16] S. Arbesman, J. M. Kleinberg, and S. H. Strogatz, "Superlinear scaling for innovation in cities," *Phys. Rev. E*, vol. 79, p. 016115, Jan 2009. [Online]. Available: <http://link.aps.org/doi/10.1103/PhysRevE.79.016115>
- [17] C. Manning, P. Raghavan, and H. Schütze, *Introduction to information retrieval*. Cambridge University Press Cambridge, 2008, vol. 1.
- [18] F. Ferreira and J. Waldfogel, "Pop internationalism: Has a half century of world music trade displaced local culture?" National Bureau of Economic Research, Tech. Rep., 2010.
- [19] D. Müllner, "Modern hierarchical, agglomerative clustering algorithms," *arXiv*, p. 29, Sep. 2011. [Online]. Available: <http://arxiv.org/abs/1109.2378>
- [20] L. Page, S. Brin, R. Motwani, and T. Winograd, "The pagerank citation ranking: Bringing order to the web." 1999.
- [21] U. Brandes and D. Wagner, "Analysis and visualization of social networks," *Graph drawing software*, pp. 321–340, 2004.